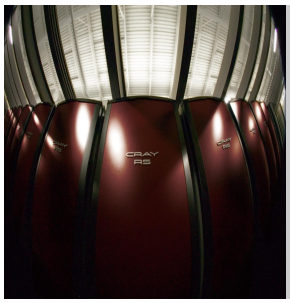

Sandia National Laboratories

Exceptional service in the national interest

D^2T: Doubly Distributed Transactions for High Performance and Distributed Computing

Jai Dayal, **Jay Lofstead**,

Karsten Schwan, Ron Oldfield

Georgia Tech, College of Computing, Atlanta, GA, USA &
Sandia National Laboratories
Scalable System Software
Albuquerque, NM, USA
gflofst@sandia.gov

HPDC 2013
June 18, 2013

U.S. DEPARTMENT OF ENERGY

NNSA National Nuclear Security Administration

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Sandia National Laboratories

Doubly Distributed Transactions for HPDC

- Need guarantees operations are complete and correct
  - Eventual consistency not good enough (space/time constraints)
  - Paxos/Zookeeper (and others) 1xN only
- Full ACID properties possible (with sufficient hardware support)
- MxN at extreme scale hard
  - 10 million clients to 10000 servers

Improvements in This Version

Sandia National Laboratories

Old Protocol	New Protocol
20M + 12N + 12a	13M + 0N + 2Na

M = Number of Clients
N = Number of Servers
a = Messages across

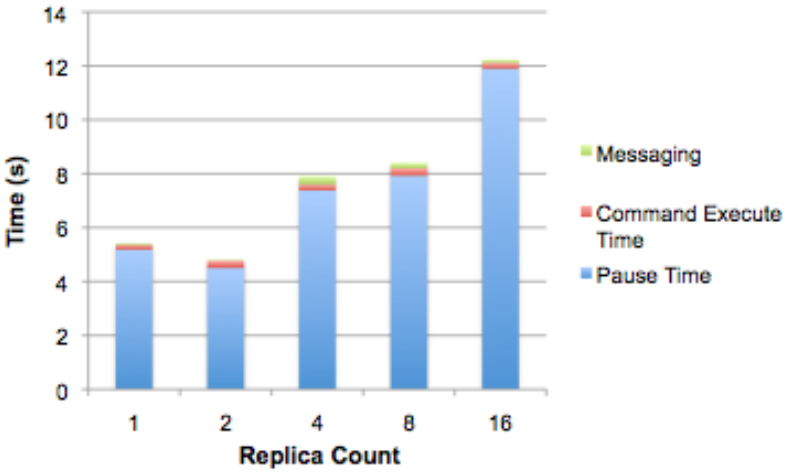

- Optimized implementation complete
- Much better scalability
- Example Services
  - Data Storage
  - Metadata
- Time spent executing transaction protocol is negligible